

EXAMINER'S AMENDMENT

1. An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in at least a telephone interview with Matt Hindman on 6/24/08.

2. Please replace Claims 1, 8, 15, 20, 27, 31, 33, 40, and 47 with currently amended Claims 1, 8, 15, 20, 27, 31, 33, 40, and 47.

Claim 1 (Currently Amended) A method of garbage collecting in a storage device comprising:

locating blocks of data stored in a contiguous circular_log that are both referenced by at least one other block of data stored within a set of one or more storage trees and within a range at a tail of the log using pruned walking, the range representing an address range within an allocated segment of the log, each storage tree representing active snapshots of user-generated data stored from each one of a plurality of remote computer systems taken at different points in time and having a plurality of nodes, each node representing a block of data of an active snapshot associated with each storage tree, the plurality of nodes including leaf nodes representing data blocks of the user-generated data stored from each one of the plurality of remote computer systems and

root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots, each block of data stored within the circular log in a sequential order, wrapping around to a beginning of the log once an end of the log is reached and the beginning of the log has been cleaned, so that the size of the contiguous log does not increase beyond its originally allocated size;

copying the blocks of data that are referenced by one or more other blocks of data of other nodes and within the range to an unallocated segment of the log, wherein blocks of data that are not referenced by other blocks of data and within the range remain untouched, and

marking the range at the tail of the log as unallocated so that at least a portion of an address space within the range ~~can be~~is reclaimed.

Claim 8 (Currently Amended) A method comprising:

garbage collecting in a circular log within a range of addresses in a storage system having a plurality of storage trees representing active snapshots of user-generated data stored from each one of a plurality of remote computer systems taken at different points in time, each storage tree having a plurality of nodes, each node representing a block of data of an active snapshot associated with each storage tree, the plurality of nodes including leaf nodes representing data blocks of the user-

generated data stored from each one of the plurality of remote computer systems and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots, and each storage tree having multiple references to a same block of data, wherein each block of data is stored within the circular log as contiguous data in a sequential order, wrapping around to a beginning of the log once an end of the log is reached and the beginning of the log has been cleaned, so that the size of the log does not increase beyond its originally allocated size, the garbage collecting including:

pruning walking of the plurality of storage trees to determine active blocks of data within said range, where active blocks of data are those that are referenced by at least one other block of data still stored within one of the plurality of storage trees, the pruning walking including:

determining, based on accessing in one of said plurality of storage trees a parent node that has a plurality of descendent nodes, that none of the plurality of descendant nodes are associated with blocks of data within the range; and

skipping the walking of the plurality of descendent nodes based on said determining,

wherein the active blocks determined to be in the range are copied out of the range and the range is marked as unallocated so that at least a portion of the address space within the range ~~can be is~~ reclaimed.

Claim 15 (Currently Amended) A method of garbage collecting in a storage system comprising:

performing the following operations in a circular log until each block of data that is active in a range to be cleaned at a tail of the log of data is copied to a head of the log, each block of data stored within the circular log as contiguous data in a sequential order, wrapping around to a beginning of the log once an end of the log is reached and the beginning of the log has been cleaned, so that the size of the log does not increase beyond its originally allocated size, wherein the range to be cleaned is a range of addresses in the storage system, the storage system having a plurality of storage trees representing active snapshots of user-generated data stored from each one of a plurality of remote computer systems taken at different points in time, each storage tree having a plurality of nodes, wherein each block of data is represented by a node of the storage tree, the plurality of nodes including leaf nodes representing data blocks of the user-generated data stored from each one of the plurality of remote computer systems and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and

sharing of the user-generated data between and among different snapshots of the plurality of snapshots, the operations including:

locating blocks of data in the log that are both referenced by at least one other block of data stored within the nodes of the plurality of storage trees and within the range to be cleaned using pruned walking;

copying blocks of data associated with child nodes of a current node that are within the range to be cleaned to the head of the log;

retrieving a block of data associated with the current node, upon determining that a minimum address value among addresses of descendent nodes is within the range to be cleaned;

designating, as the current node, one of the child nodes of the current node that is an interior node, upon determining that at least one child node is an interior node;

designating, as the current node, an ancestor node of the current node whose descendent nodes are unprocessed; and

marking the range as unallocated when the blocks of data that are active and within the range are copied to the head of the log so that at least a portion of the address space within the range to be cleaned ~~can be is~~ reclaimed.

Claim 20 (Currently Amended) A system comprising:

a storage device to store a number of blocks of data in a circular log of contiguous data, wherein the blocks of data that are marked as allocated are non-

modifiable, wherein each block of data within the circular log is stored in a sequential order, wrapping around to a beginning of the log once an end of the log is reached and the beginning of the log has been cleaned, so that the size of the log does not increase beyond its originally allocated size; and

 a garbage collection logic to locate the blocks of data that are both referenced by at least one other block stored within a set of one or more storage trees, and within a range at a tail of the log using pruned walking, the range representing an address range to be cleaned within an allocated segment of the log,

 wherein the log is implemented in a hierarchical architecture having a plurality of storage trees, each representing a snapshot taken at a point in time of target data being processed, each storage tree having a plurality of nodes, and each node representing a block of data of an active snapshot associated with each storage tree, each of the active snapshots including user-generated data stored from a file system located on a remote computer taken at a different time, each storage tree having a plurality of nodes including leaf nodes representing data blocks of the user-generated data stored from the remote computer file system and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots,

and wherein said garbage collection logic is operable to copy the located blocks of data to a head of the log and mark the range to be cleaned as unallocated so that at least a portion of the address space within the range ~~can be~~is reclaimed.

Claim 27 (Currently Amended) A backup system comprising:

a plurality of storage trees stored in a contiguous circular log, each storage tree representing an active snapshot of user-generated data stored from a file system located on a remote computer taken at a different time, each storage tree having a plurality of nodes, including leaf nodes representing data blocks of the user-generated data stored from the remote computer file system and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated ~~target~~-data between and among different snapshots of the plurality of snapshots, wherein each leaf node of user-generated data from said remote computer file system has been backed up from a set of one or more storage devices within the remote computer, each block of data is stored within the circular log in a sequential order, wrapping around to a beginning of the log once an end of the log is reached and the beginning of the log has been cleaned, so that the size of the contiguous log does not increase beyond its originally allocated size;

a storage space to store said blocks of data having been allocated from a backup storage space in said set of one or more storage devices;

a set of one or more location tables having stored therein a minimum address value for descendent nodes of interior nodes of said plurality of storage trees; and a garbage collection logic to clean a currently selected range from the tail of said log, said garbage collection logic to perform pruned walking of nodes of said plurality of storage trees based on said set of location tables and said currently selected range to locate blocks of data that are referenced by at least one other block of data currently stored within the plurality of storage trees.

Claim 31 (Currently Amended) An apparatus comprising:

a hardware backup system to backup a file system located on a remote computer by storing the file system to a circular log as contiguous blocks of data, said backup file system including:
a tracking logic to generate a plurality of storage trees, each storage tree representing active backup snapshots of user-generated data stored from said remote computer file system taken at different times by recording references to blocks of backed up data stored in a set of one or more storage devices, each storage tree having a plurality of nodes including leaf nodes representing data blocks of the user-generated data stored from the remote computer file system and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated ~~target~~ data between and among different

snapshots of the plurality of snapshots, each block of data stored within the circular log in a sequential order, wrapping around to a beginning of the log once an end of the log is reached and the beginning of the log has been cleaned, so that the size of the contiguous log does not increase beyond its originally allocated size;

an allocator logic to allocate contiguous blocks of storage space from a log of a back up storage space to store said blocks of backed up data;

a garbage collection logic to, responsive to deletion of one or more of said active backup snapshots, clean a currently selected contiguous range from the tail of said log, said garbage collection logic to,

walk only those nodes of said plurality of storage trees that possibly identify those of said blocks of data that are stored in said currently selected contiguous range or that possibly are themselves stored in said currently selected contiguous range,

locate blocks of data that are referenced by at least one other block of data stored within the plurality of storage trees, and

sweep said currently selected contiguous range, copying blocks of data that are referenced and within the range out of the range and marking the range as unallocated so that at least a portion of the address space within the range ~~can be~~is reclaimed.

Claim 33 (Currently Amended) A computer-storage medium that stores instructions, which when executed by a computer, cause said computer to perform garbage collection operations in a hardware storage device comprising:

locating blocks of data stored in a circular log that are both referenced by at least one other block of data stored within a plurality of storage trees and within a range at a tail of the log using pruned walking, the range representing an address range to be cleaned within an allocated segment of the log, each storage tree representing active snapshots of user-generated data stored from each one of a plurality of remote computer systems at different points in time and having a plurality of nodes, and each node representing a block of data of an active snapshot associated with each storage tree, the plurality of nodes including leaf nodes representing data blocks of the user-generated data stored from each one of the plurality of remote computer systems and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots, wherein each block of data within the circular log is stored as contiguous data in a sequential order, wrapping around to a beginning of the log once an end of the log is reached and the beginning of the log has been cleaned, so that the size of the log does not increase beyond its originally created size;

copying the blocks of data that are referenced and within the range to an unallocated segment of the log, wherein blocks of data that are not referenced by other blocks of data and within the range remain untouched, and

marking the range at the tail of the log as unallocated so that at least a portion of an address space within the range ~~can be~~is reclaimed.

Claim 40 (Currently Amended) A computer-storage-medium that stores instructions, which when executed by a computer, cause said computer to perform operations in a hardware storage device comprising:

garbage collecting within a range of addresses in circular log of contiguous blocks of data having a plurality of storage trees stored therein, each storage tree representing active snapshots of user-generated data stored from each one of a plurality of remote computer systems taken at different times and having a plurality of nodes, the plurality of nodes including leaf nodes representing data blocks of the user-generated data stored from each one of the plurality of remote computer systems and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated ~~target~~ data between and among different snapshots of the plurality of snapshots and having multiple references to the same block of data, each block of data stored within the circular log in a sequential order, wrapping around to a beginning of the log once an end of the log is reached and the beginning of the log has been

cleaned, so that the size of the contiguous log does not increase beyond its originally allocated size, the garbage collecting including:

pruning walking of the plurality of storage trees to determine active blocks of data within said range, where active blocks of data are those that are referenced by at least one other block of data still in one of the plurality of storage trees, the pruning walking including:

determining, based on accessing in one of said plurality of storage trees a parent node that has a plurality of descendent nodes, that none of the plurality of descendent nodes are associated with blocks of data within the range; and

skipping the walking of the plurality of descendent nodes based on said determining,

wherein the active blocks determined to be in the range are copied out of the range and the range is marked as unallocated so that at least a portion of the address space within the range ~~can be is~~ reclaimed.

Claim 47 (Currently Amended) A computer-storage medium that stores instructions, which when executed by a computer, cause said computer to perform garbage collection operations in a hardware storage device comprising:

performing the following operations in a circular log of contiguous data blocks until each block of data that is active in a range to be cleaned at a tail of the log of contiguous data is copied to a head of the log, wherein the range to be cleaned is a

range of addresses in a storage system having a plurality of storage trees, each storage tree having a plurality of nodes, wherein a block of data is associated with a node of the storage tree, each node representing a block of data of an active snapshot associated with each storage tree, each of the active snapshots including user-generated data stored from a file system located on a remote computer taken at a different time, each storage tree having a plurality of nodes including leaf nodes representing data blocks of the user-generated data stored from the remote computer file system and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots, and wherein data stored within the log is stored in a sequential order, wrapping around to a beginning of the log once an end of the log is reached and the beginning of the log is cleaned, so that the size of the log does not increase beyond its originally allocated size, the operations including:

locating blocks of data in a log that are both referenced by at least one other block of data stored within the plurality of storage trees, and within the range using pruned walking;

copying blocks of data associated with child nodes of a current node that are within the range to be cleaned to the head of the log;

retrieving a block of data associated with the current node, upon determining that a minimum address value among addresses of descendent nodes is within the range to be cleaned;

designating, as the current node, one of the child nodes of the current node that is an interior node, upon determining that at least one child node is an interior node;

designating, as the current node, an ancestor node of the current node whose descendent nodes are unprocessed; and

marking the range at the tail of the log as unallocated when the blocks of data that are active and within the range are copied to the head of the log so that at least a portion of an address space within the range ~~can be~~is reclaimed.

Allowable Subject Matter

3. Claims 1, 3-11, 13-18, 20-22, 24-33, 35-43, and 45-50 (renumbered 1-44) are allowed.
4. The following is an examiner's statement of reasons for allowance:

Prior art(s) do not appear to anticipate or make obvious at least the features of:

garbage collecting in a circular log;

storage trees of snapshots where system generated data includes pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots; and

the pruned walking method for garbage collecting in the storage trees/circular log.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

Interview Summary

5. An examiner-initiated interview was held on at least 6/24/08. During the interview, claim language was agreed upon to issue the instant application. Specifically, it was agreed to amend the claims via an examiner's amendment to have the claims include the user-generated and system-generated data and amend around various new 112 rejections discovered.

Conclusion

6. Any prior art made of record and not relied upon is considered pertinent to applicant's disclosure. Applicant is advised that, although not used in the rejections above, prior art cited on any PTO-892 form and not relied upon is considered materially relevant to the applicant's claimed invention and/or portions of the claimed invention.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Brent S. Stace whose telephone number is 571-272-8372 and fax number is 571-273-8372. The examiner can normally be reached on M-F 9am-5:30pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Apu M. Mofiz can be reached on 571-272-4080. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/BRENT STACE/
Examiner, Art Unit 2161

/Apu M Mofiz/

Supervisory Patent Examiner, Art Unit 2161